

SEMICONDUCTORS / PROCESSORS

FEATURE

Transistor Aging

Measuring the degradation of microprocessors is tricky. Doing it better would unleash more processing power

By JOHN KEANE, CHRIS H. KIM / MAY 2011




Illustration: Daniel Hertzberg

You know when it's time to get a new car. Your odometer is far into six digits, perhaps the engine is burning lots of oil, or the transmission is growling. Fixing all that might well cost quite a bit more than your ancient vehicle is worth.

But what about your microprocessor? Unlike automobiles, microprocessors don't have convenient little gauges that reflect how much wear and tear they've endured. And wear they do—though you'll probably never notice it. The degradation of their transistors over time leads slowly but surely to decreased switching speeds, and it can even result in outright circuit failures.

You generally don't perceive this deterioration because semiconductor companies always play it very safe—they set the clock-speed rating of their microprocessors so conservatively that almost every one of their products will continue to operate flawlessly throughout its intended lifetime. That strategy works. But it's kind of like never taking your Ferrari out of the slow lane because you're concerned that its engine might throw a rod 10 years down the road.

Several different phenomena can degrade the transistors on chips. How those phenomena combine to diminish a chip's functioning depends on such factors as the circuit arrangement of the aging transistors as well as the voltages and temperatures they're exposed to. With all these variables, it's difficult to predict how the peak performance of a given microprocessor will decline over time.

We and other researchers are trying to improve that situation. One critical aspect of the [work we did](#) at the University of Minnesota was to develop better ways to study the different physical mechanisms of transistor aging. Today semiconductor engineers measure those aging effects primarily by examining transistors one at a time, using microscopic electrodes to probe a silicon wafer. The necessary equipment can cost tens of millions of dollars, and probing transistors individually is arduous when you're trying to gather many observations. Sometimes you can't do those measurements well, no matter how much time you spend.

We need better techniques. And we need them sooner rather than later. Microprocessors now contain billions of transistors, sometimes operating at clock speeds in excess of 3 billion cycles per second. The blazingly fast clocks mean that the transistors are exposed to lots of heat, which accelerates their decline. Another worry is that there are precariously small voltage differences between supply levels and the threshold at which the transistors turn on. Also, various improvements in the way silicon logic is fabricated have introduced new concerns about degradation. And transistors scaled down to today's tiny dimensions experience more variation than ever in their operating conditions, which in turn leads to great differences from one transistor to another in how fast they wear out.

With better ways to measure transistor aging, chipmakers could let their microprocessors run faster—appreciably faster—than they do now. In the future it might even be possible to use these techniques to build circuits into microprocessors that continuously measure the subtle effects of aging and adjust clock frequency or operating voltages so that the transistors, old or new, could always run at peak speeds.




Photo: Cascade Microtech

SQUINT HARD:

The traditional method of monitoring transistor-aging effects in chips requires the careful placement of tiny probes, which are manipulated while being viewed under a microscope.

INFOGRAPHIC: Silicon's Decline



Why should transistors age at all? The kinds of transistors we're talking about here—the good old metal-oxide semiconductor field-effect transistors that are the basis of ordinary CMOS chips—

function as electrical switches. A MOSFET has four terminals, called the body, the gate, the source, and the drain, although the source and body are often connected. The voltage that's applied to the gate determines whether current can flow between the source and drain. Although a thin layer of dielectric material electrically insulates the gate, the electric field applied across it alters the conductivity of the underlying semiconductor channel connecting the source and drain.

And that brings us to our first degenerative mechanism: Over time, charge carriers (electrons for negative, or *n*-channel, MOSFETs; holes for positive, or *p*-channel, MOSFETs) with a little more energy than the average will stray out of the conductive channel between the source and drain and get trapped in the insulating dielectric. This process, called hot-carrier injection, eventually builds up electric charge within the dielectric layer, increasing the voltage needed to turn the transistor on. As this threshold voltage increases, the transistor switches more and more slowly.

There's a second mechanism that can also trap charge in the dielectric, and it doesn't require any current to flow between the source and drain. Whenever you apply voltage to the gate, a phenomenon called bias temperature instability can cause a buildup of charge in the dielectric, along with other subtle problems. After that gate voltage is removed, though, some of this effect spontaneously disappears. This recovery occurs within a few tens of microseconds, making it difficult to observe during routine experiments, where you stress the transistor but measure the resulting effects only after the stress is removed.

Yet another aging mechanism comes into play when a voltage applied to the gate creates electrically active defects, known as traps, within the dielectric. If they become too numerous, these charge traps can join and form an outright short circuit between the gate and the current channel. This kind of failure is called oxide breakdown, or more verbosely, time-dependent dielectric breakdown. Unlike the other aging mechanisms, which cause a gradual decline in performance, the breakdown of the dielectric can lead to the catastrophic failure of the transistor, causing the circuit it's in to malfunction.

As if the aging of transistors wasn't enough to worry about, semiconductor engineers also have to grapple with the metal connections between transistors wearing out over time. The concern here is a phenomenon called [electromigration](#), which damages the copper or aluminum connections that tie transistors together or link them to the outside world.

Electromigration occurs when a surge of current knocks metal atoms loose and causes them to drift along with the flow of electrons. This depletes the metal of some of its atoms upstream, while causing a buildup of metal downstream. The

upstream thinning of the metal increases the resistance of the connection, sometimes to the point that it can become an open circuit. While downstream deposition isn't similarly catastrophic, it can cause the metal to bulge out of its designated track.

Despite the many efforts of process engineers to create long-lasting transistors and metallic connections, a certain amount of wear is unavoidable. So it must be reckoned with. Probing chips to study how their transistors degrade provides only limited information on how bad the problems are inside a real microprocessor. A better experimental approach, just now taking hold in industry, is to fabricate special chips for the sole purpose of testing how the transistors on them function over time. These chips contain what are known as ring oscillators, consisting of many inverter circuits chained together in a loop. Each inverter outputs the opposite of whatever signal is applied to its input. So when an odd number of them are wired into a ring, the circuit oscillates, at a frequency that depends on how fast the constituent transistors can switch states.

To measure the slowdown that arises as transistors age, engineers subject chips containing such ring oscillators to extreme conditions, raising the supply voltage or operating temperature so that they wear out in a matter of hours or days. Engineers can measure how transistor-switching times increase, or how the average time to failure decreases, during several different accelerated-aging experiments conducted using a variety of stress levels. That allows them to extrapolate their results to the real world, where transistors age much more slowly than they do during these tests.





Illustration: Daniel Hertzberg



Testing ring oscillators provides more insight into how circuits age than does probing transistors individually, but it still has flaws. One is that it takes a relatively long time to make a measurement—a large fraction of a second. That's because you're looking for very subtle changes in the frequency of the oscillations, so you need to count a lot of cycles to measure the shift.

Taking a second might seem quick enough, but remember, most of one particular aging effect (bias temperature instability) lasts for only a few tens of *microseconds* after the stress is removed.

Most often, you would stress the circuit by raising its supply voltage; elevating the temperature works less well. But you then need to return the voltage to its normal level to judge how much the transistors have changed. So most of the effects of bias temperature instability will disappear before you'll have a chance to observe them. We've worked hard to remedy this problem and also to improve the way transistors are examined in general.

To measure the shifting frequency of ring oscillators much more quickly, we developed what we call a silicon odometer. It is based on a pair of ring oscillators and works by measuring their beat frequency. If you're at all musical, you're probably familiar with the beat phenomenon: When two very similar notes are played simultaneously, you hear just one

note whose amplitude changes rhythmically—it beats. The number of beats per second is equal to the difference in frequency between the two original notes.

In our silicon odometer, we measure the difference between the frequency of an unstressed ring oscillator and the one we have stressed by increasing its operating voltage. To measure actual beats, we'd have to sum the two outputs and run the results through an analog-to-digital converter. Rather than doing that, we do something similar that is easier to accomplish: We sample the output of one oscillator at intervals set by the output of the other. That gives us a digital signal that oscillates at the beat frequency, which is easy to measure with simple circuitry. This approach allows us to measure changes in transistor switching times as small as one part in 10 000 in less than a microsecond, which is short enough to capture even the most fleeting effects of aging.

In addition to the short measurement times, our silicon odometer has another nice attribute: It's essentially immune to the gradual changes in operating voltage or temperature that often take place during long stress experiments. Any stray variations in either will shift the frequencies of the two oscillators by roughly the same amount. So their difference, which is all we measure, will continue to reflect just the effects of stress. We've gathered a lot of valuable aging data with this odometer. And with a somewhat different design, which we've dubbed the all-in-one odometer, we were able to separate the simultaneous effects of hot-carrier injection from those of bias temperature instability.

To get a better idea of how aging varies from place to place on a chip, we developed a third design, a statistical silicon odometer, which uses arrays of stressed ring oscillators. The circuitry we've devised allows us to connect any of those ring oscillators to a reference oscillator on the chip that runs at about the same frequency. We can thus measure subtle aging effects at many different locations on the silicon die.

Beat-frequency detection systems are not the only on-chip monitors we have implemented, though. We've put together other circuits to study the statistics of oxide breakdown, which is rarely extensive enough to disable a transistor. But rare events count when you're dealing with the billions of transistors in a modern microprocessor, which make for billions of opportunities for failure. To study that, we have designed several different arrays of transistors that are particularly prone to having their gate dielectrics wear out and finally break down. By stressing all the transistors in these arrays at the same time, we can easily take measurements of their time to breakdown. This information is useful to chip designers, whose usual recourse is to deal with this problem by adjusting fabrication procedures.

Our on-chip circuit-aging sensors are providing valuable information about how real circuits age. Not long from now, we expect it will be possible to design similar circuits into microprocessors for use as real-time aging monitors, which could trigger adjustments in clock speeds or operating voltages to ensure that these chips work at peak levels, even as they grow older. That would eliminate the need for slowing the chips down from the start.

The semiconductor industry is just beginning to dabble with such ideas. Intel, for example, has engineered a technology it code-named Foxtan, which measures the amount of propagation delay along certain critical signal paths within a microprocessor. Specialized circuitry on the microprocessor then uses that information to adjust supply voltage and clock frequency. Intel had planned to include Foxtan technology in its dual-core Itanium 2 processor, released in 2006, but undisclosed complications delayed the introduction. We suspect that using something like our silicon odometer might help overcome the difficulties Intel faced.

In that vein, we are collaborating with groups at Intel and also at such companies as Advanced Micro Devices, Broadcom, Freescale Semiconductor, Texas Instruments, and (with help from the Defense Advanced Research Projects Agency) IBM, which are all striving to eliminate the compromises they routinely make to accommodate transistor aging. Silicon odometers, or something similar, might well appear in some of their new CMOS chips.

With such technology, microprocessors could be made to always perform at their peak levels, even though those levels would still decline slowly over time—as is the case with so many of our other machines. This will take some getting used to, of course, particularly if you ever attempt to sell your used computer. Maybe sometime in the not-too-distant future we'll be reading classified ads that say things like, "For sale: Five-year-old laptop, but with like-new processor. Used only to compute at church on Sundays."

This article originally appeared in print as "An Odometer for CPUs."

About the Author

John Keane and Chris H. Kim met in 2005, when Keane, then a graduate student (and now a component-design engineer for Intel in Hillsboro, Ore.), joined Kim's research group at the University of Minnesota. To study the wear and tear transistors experience, Keane and Kim designed special-purpose chips. Making your own integrated circuits is tricky and expensive, but it's an addictively productive way to do this kind of research, Kim says: "Once you start building chips, it's hard to go back."

[Recommend](#)

16 recommendations. Sign Up to see what your friends recommend.